

Notice of Allowability	Application No.	Applicant(s)	
	10/767,182	PILLAI ET AL.	
	Examiner	Art Unit	
	HIEU HOANG	2452	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address--

All claims being allowable, PROSECUTION ON THE MERITS IS (OR REMAINS) CLOSED in this application. If not included herewith (or previously mailed), a Notice of Allowance (PTOL-85) or other appropriate communication will be mailed in due course. **THIS NOTICE OF ALLOWABILITY IS NOT A GRANT OF PATENT RIGHTS.** This application is subject to withdrawal from issue at the initiative of the Office or upon petition by the applicant. See 37 CFR 1.313 and MPEP 1308.

1. ☒ This communication is responsive to communication on 09/06/2011.
2. ☐ An election was made by the applicant in response to a restriction requirement set forth during the interview on ____; the restriction requirement and election have been incorporated into this action.
3. ☒ The allowed claim(s) is/are 28-43.
4. ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 - a) ☐ All b) ☐ Some* c) ☐ None of the:
 1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. ____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this national stage application from the International Bureau (PCT Rule 17.2(a)).

* Certified copies not received: ____.

Applicant has THREE MONTHS FROM THE "MAILING DATE" of this communication to file a reply complying with the requirements noted below. Failure to timely comply will result in ABANDONMENT of this application.

THIS THREE-MONTH PERIOD IS NOT EXTENDABLE.

5. ☐ A SUBSTITUTE OATH OR DECLARATION must be submitted. Note the attached EXAMINER'S AMENDMENT or NOTICE OF INFORMAL PATENT APPLICATION (PTO-152) which gives reason(s) why the oath or declaration is deficient.
 6. ☐ CORRECTED DRAWINGS (as "replacement sheets") must be submitted.
 - (a) ☐ including changes required by the Notice of Draftsperson's Patent Drawing Review (PTO-948) attached
 - 1) ☐ hereto or 2) ☐ to Paper No./Mail Date ____.
 - (b) ☐ including changes required by the attached Examiner's Amendment / Comment or in the Office action of Paper No./Mail Date ____.
- Identifying indicia such as the application number (see 37 CFR 1.84(c)) should be written on the drawings in the front (not the back) of each sheet. Replacement sheet(s) should be labeled as such in the header according to 37 CFR 1.121(d).**
7. ☐ DEPOSIT OF and/or INFORMATION about the deposit of BIOLOGICAL MATERIAL must be submitted. Note the attached Examiner's comment regarding REQUIREMENT FOR THE DEPOSIT OF BIOLOGICAL MATERIAL.

Attachment(s)

- | | |
|--|---|
| 1. <input checked="" type="checkbox"/> Notice of References Cited (PTO-892) | 5. <input type="checkbox"/> Notice of Informal Patent Application |
| 2. <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | 6. <input checked="" type="checkbox"/> Interview Summary (PTO-413),
Paper No./Mail Date <u>9/20/2011</u> . |
| 3. <input type="checkbox"/> Information Disclosure Statements (PTO/SB/08),
Paper No./Mail Date ____ | 7. <input checked="" type="checkbox"/> Examiner's Amendment/Comment |
| 4. <input type="checkbox"/> Examiner's Comment Regarding Requirement for Deposit
of Biological Material | 8. <input checked="" type="checkbox"/> Examiner's Statement of Reasons for Allowance |
| | 9. <input type="checkbox"/> Other ____. |

/HIEU HOANG/
Primary Examiner, Art Unit 2452

DETAILED ACTION

This office action is in response to the communication filed on 09/06/2011.

EXAMINER'S AMENDMENT

An examiner's amendment to the record appears below. Should the changes and/or additions be unacceptable to applicant, an amendment may be filed as provided by 37 CFR 1.312. To ensure consideration of such an amendment, it MUST be submitted no later than the payment of the issue fee.

Authorization for this examiner's amendment was given in a telephone interview with on Jeremy Curcuri on 09/20/2011.

The claims have been amended as follows:

1 - 27. (Canceled)

28. (Currently Amended) A method for performing remote access commands between nodes, the method comprising:

establishing a mutually pre-agreed upon data allotment for a first node and a second node;

detecting an application request in a request queue of the first node, the application request identifying a data access task to be performed between the first node and the second node, wherein the application request is a remote direct memory access (RDMA) request for the first node to access data in a memory at the second

Art Unit: 2452

node and wherein the first node and the second node exchange requests and responses to perform the data access task via a first channel adapter at the first node and a second channel adapter at the second node; and

assigning a context including a set of channel adapter resources of the first node to process the application request to completion of the data access task, the context operating a task manager that tracks an offset for the data access task, wherein the set of channel adapter resources are hardware resources within the first channel adapter of the first node;

issuing a first request from the first node to the second node, the first request requesting the data access task be performed between the first node and the second node;

receiving, at the first node, a first response from the second node that partially completes the data access task and that contains ~~data in an~~ a first amount of data not exceeding the data allotment;

issuing at least one subsidiary request from the first node to the second node to further complete the data access task between the first node and the second node, the at least one subsidiary request based on an amount of partial completion of the data access task between the first node and the second node; and

receiving, from the second node in response to the at least one subsidiary request, at least one ~~corresponding~~ subsidiary response that further completes the data access task between the first node and the second node and that contains ~~data in an~~ a second amount of data not exceeding the data allotment;

Art Unit: 2452

wherein a new starting value for subsidiary data to be transferred for each of the
at least one subsidiary request is calculated from the offset; and

wherein the context is ~~pre-emptible~~ pre-empted between the first response and a
first subsidiary request, and between a subsidiary response and a subsequent
subsidiary request, wherein the pre-empted context is available to be assigned to
~~support~~ other data access tasks by the first channel adapter performing a context switch
from the data access task to another data access task of the other data access tasks.

29. (Currently Amended) The method of claim 28₁ further comprising:

pre-empting the context prior to ~~full~~ the completion of the data access task;

issuing a second request from the first node to the second node for ~~[[an]]~~ the
other data access task; and

resuming the context for the other data access task.

30. (Previously Presented) The method of claim 29₁ wherein the other data access task
is identified by a different application request in a different request queue.

31. (Currently Amended) The method of claim 28₁ wherein the step of issuing the at
least one subsidiary request comprises:

calculating a remaining amount of data required to complete the data access task
between the first node and the second node; and

Art Unit: 2452

creating ~~[[a]]~~ the at least one subsidiary request ~~to reference at least a portion of~~
for the remaining amount of data required to complete the data access task.

32. (Currently Amended) The method of claim 31, wherein the step of calculating the remaining amount of data comprises:

determining a total completed amount of data processed for the data access task by the first request and ~~associated~~ the first response and ~~all~~ the at least one subsidiary request~~[[s]]~~ and ~~corresponding~~ the at least one subsidiary response~~[[s]]~~ between the first node and the second node; and

determining the remaining amount of data required to complete the data access task as a difference between an initial amount of data specified by ~~[[an]]~~ the application request and the total completed amount of data.

33. (Currently Amended) The method of claim 28, wherein:

~~the first and second nodes are nodes that utilize channel adapters to exchange the first request and the at least one subsidiary request and the corresponding first response and the at least one subsidiary response;~~

~~the application request is a remote direct memory access request for the first node to access data in a memory at the second node; and~~

~~[[an]]~~ the initial amount of data specified by the application request is a total amount of data that the first node is to access in the memory at the second node.

Art Unit: 2452

34. (Currently Amended) The method of claim 33, wherein the first request and the at least one subsidiary request are read ~~remote direct memory access~~ RDMA commands issued by the first node to ~~read data in the memory from~~ the second node.

35. (Currently Amended) The method of claim 28, wherein the step of establishing the data allotment comprises:

dynamically determining the data allotment between the first node and the second node[[s]] based on at least one external data allotment event, ~~such that~~ wherein if the at least one external data allotment event occurs, the first node and the second node[[s]] change ~~a value of~~ the data allotment.

36. (Currently Amended) A ~~computerized~~ first device including a ~~communications interface~~ channel adapter, the ~~communications interface~~ channel adapter comprising:

a processor within the ~~communications interface~~ channel adapter;
a memory coupled to the processor; and
a communications port coupled to the processor and the memory; and
~~an interconnection mechanism coupling the processor and the communications~~
~~port~~;

wherein the processor executes ~~logic of a communications interface application~~ channel adapter codes stored in the memory to ~~form a communications interface~~ process that perform[[s]] remote access commands between nodes the first device and a second device by performing the ~~operations~~ steps of:

Art Unit: 2452

establishing a mutually pre-agreed upon data allotment for ~~[[a]]~~ the first node device and ~~[[a]]~~ the second node device;

detecting an application request in a request queue of the first device, the application request identifying a data access task to be performed between the first device and the second node device, wherein the application request is a remote direct memory access (RDMA) request for the first device to access data in another memory at the second device and wherein the first device and the second device exchange requests and responses to perform the data access task via the channel adapter of the first device and another channel adapter at the second device; and

assigning a context including a set of channel adapter resources of the first ~~node device~~ device to process the application request to completion of the data access task, the context operating a task manager that tracks an offset for the data access task, wherein the set of channel adapter resources are hardware resources within the channel adapter of the first device;

issuing a first request from the first ~~node device~~ device to the second ~~node device~~ device, the first request requesting the data access task be performed between the first ~~node device~~ device and the second ~~node device~~ device;

receiving, at the first ~~node device~~ device, a first response from the second ~~node device~~ device that partially completes the data access task and that contains ~~data in an~~ a first amount of data not exceeding the data allotment;

issuing at least one subsidiary request from the first ~~node device~~ device to the second ~~node device~~ device to further complete the data access task between the first ~~node device~~ device and

Art Unit: 2452

the second ~~node~~ device, the at least one subsidiary request based on an amount of partial completion of the data access task between the first ~~node~~ device and the second ~~node~~ device; and

receiving, from the second ~~node~~ device in response to the at least one subsidiary request, at least one ~~corresponding~~ subsidiary response that further completes the data access task between the first ~~node~~ device and the second ~~node~~ device and that contains ~~data in an~~ a second amount of data not exceeding the data allotment;

wherein a new starting value for subsidiary data to be transferred for each of the at least one subsidiary request is calculated from the offset; and

wherein the context is ~~pre-emptible~~ pre-empted between the first response and a first subsidiary request, and between a subsidiary response and a subsequent subsidiary request, wherein the pre-empted context is available to be assigned to support other access tasks by the channel adapter performing a context switch from the data access task to another data access task of the other data access tasks.

37. (Currently Amended) The ~~computerized~~ first device of claim 36, wherein the ~~communications interface application processor~~ further executes the channel adapter codes to perform[[s]] the operations further steps of:

pre-empting the context prior to ~~full~~ the completion of the data access task;

issuing a second request from the first ~~node~~ device to the second device for ~~another~~ the other data access task; and

resuming the context for the other data access task.

38. (Previously Presented) The ~~computerized~~ first device of claim 37, wherein the other data access task is identified by a different application request in a different request queue.

39. (Currently Amended) The ~~computerized~~ first device of claim 36, wherein the step of issuing the at least one subsidiary request comprises:

calculating a remaining amount of data required to complete the data access task between the first ~~node~~ device and the second ~~node~~ device; and

creating ~~[[a]]~~ the at least one subsidiary request ~~to reference at least a portion of~~ for the remaining amount of data required to complete the data access task.

40. (Currently Amended) The ~~computerized~~ first device of claim 39, wherein the step of calculating the remaining amount of data comprises:

determining a total completed amount of data processed for the data access task by the first request and ~~associated~~ the first response and ~~all~~ the at least one subsidiary request~~[[s]]~~ and ~~corresponding~~ the at least one subsidiary response~~[[s]]~~ between the first and second ~~node~~ device; and

determining the remaining amount of data required to complete the data access task as a difference between an initial amount of data specified by ~~[[an]]~~ the application request and the total completed amount of data.

Art Unit: 2452

41. (Currently Amended) The ~~computerized~~ first device of claim 36, wherein:

~~the first and second nodes are nodes that utilize channel adapters to exchange the first request and the at least one subsidiary request and the corresponding first response and the at least one subsidiary response;~~

~~the application request is a remote direct memory access request for the first node to access data in a memory at the second node; and~~

[[an]] the initial amount of data specified by the application request is a total amount of data that the first ~~node~~ device is to access in the other memory at the second ~~node~~ device.

42. (Currently Amended) The ~~computerized~~ first device of claim 41, wherein the first request and the at least one subsidiary request are ~~read remote direct memory access~~ RDMA commands issued by the first ~~node~~ device to ~~read data in the memory from the~~ second ~~node~~ device.

43. (Currently Amended) The ~~computerized~~ first device of claim 36, wherein the step of establishing the data allotment comprises:

dynamically determining the data allotment between the first device and the second device ~~nodes~~ based on at least one external data allotment event, ~~such that~~ wherein if the at least one external data allotment event occurs, the first device and the second device ~~nodes~~ change ~~a value of~~ the data allotment.

Reasons for Allowance

The following is an examiner's statement of reasons for allowance:

Claims 1-27 are cancelled.

Claims 28-43 are allowed. The prior art of record does not teach the claimed invention, as follows.

In interpreting the currently amended claims, in light of the Specification and Applicant's arguments filed on 09/06/2011, the Examiner finds the claimed invention to be patentably distinct from the prior art of records. Specifically, the prior art of records, individually or in combination, fail to explicitly teach, suggest or render obvious the claimed invention as recited in independent claim 28 or 36.

Any comments considered necessary by applicant must be submitted no later than the payment of the issue fee and, to avoid processing delays, should preferably accompany the issue fee. Such submissions should be clearly labeled "Comments on Statement of Reasons for Allowance."

Conclusion

The prior art made of record and not relied upon is considered pertinent to applicant's disclosure:

- Pandya. US 2004/0037319, US 2004/0030757. TCP/IP processor and engine using RDMA.
- Henry et al. US 2004/0122987. Read prefetch in a storage architecture.
- Hui. US 2003/0220983. Partial downloading of objects, Infiniband.
- Blankenship. US 2003/0204679. Partial write request , Infiniband.
- Morrison et al. US 6,581,086. Task resume during a context save.
- Frink. US 6,961,801. DMA resume in a context switch.
- Philbrick et al. US 2005/0204058. Buffer management.
- Kohli et al. US 6,252,600. Dual FIFO interface.
- Gil. US 6,904,507. Buffer management in Infiniband.
- Pfister et al. US 6,832,297. Distributed buffer system.
- Fineberg. US 2004/0252709. RDMA transfer and buffer management.
- Craddock et al. US 2003/0018828. Mixed semantic Ethernet I/O path.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Hieu T. Hoang whose telephone number is 571-270-1253. The examiner can normally be reached on Monday-Thursday, 8 a.m.-5 p.m., EST.

Art Unit: 2452

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Thu Nguyen can be reached on 571-272-6967. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/Hieu Hoang/

Primary Examiner, Art Unit 2452